arXiv:1408.3651v1 [q-bio.PE] 15 Aug 2014

Is information a selectable trait?

Masoud Mirmomeni[1,2,†], William Punch[1,2,⋆], Christoph Adami[2,3,4,‡]

[1] Department of Computer Science and Engineering, Michigan State University, East Lansing, Michigan, 48824 (USA).
[2] BEACON Center for the Study of Evolution in Action, Michigan State University, East Lansing, Michigan, 48824 (USA).
[3] Department of Microbiology and Molecular Genetics, Michigan State University, East Lansing, Michigan, 48824 (USA).
[4] Department of Physics and Astronomy, Michigan State University, East Lansing, Michigan, 48824 (USA).

[†]mirmomen@msu.edu
[⋆] punch@cse.msu.edu
[‡] adami@msu.edu (corresponding author)

Abstract

There is little doubt in scientific circles that–counting from the origin of life towards today–evolution has led to an increase in the amount of information stored within the genomes of the biosphere. This trend of increasing information on average likely holds for every successful line of descent, but it is not clear whether this increase is due to a general law, or whether it is a secondary effect linked to an overall increase in fitness. Here, we use "digital life" evolution experiments to study whether information is under selection if treated as an organismal trait, using the Price equation. By measuring both sides of the equation individually in an adapting population, the strength of selection on a trait appears as a "gap" between the two terms of the right-hand-side of the Price equation. We find that information is strongly selected (as it encodes all fitness-producing traits) by comparing the strength of selection on information to a weakly selected trait (sequence length), as well as to a neutral marker. We observe that while strength of selection on arbitrary traits can vary during an experiment (including reversing sign), information is a selectable trait that must increase in a fixed environment.

Whether or not complexity increases in biological evolution is a hotly debated topic. The controversy is, of course, intensified by public uneasiness with evolutionary concepts. However, even within firmly scientific circles the question is considered unsettled for a variety of reasons. While no agreed-upon definitions of complexity exist (Adami, 2002), it is tempting to use information as a proxy for complexity, perhaps based on the vague notion that "it takes information do produce something complex". However, a precise definition of "information" in the context of genomes is thought to be lacking. Furthermore, historical evidence of periods of extinction suggests that information may not always have been increasing. But even if it was stipulated that information increases, it is not clear whether this trend is active or passive, in other words, whether information is actively selected for in the evolutionary process, or whether it increases simply because it is correlated with another trait that is under selection.

While it is clear and obvious to every researcher in the field of evolutionary biology (and most often beyond) that information is "that which resides in the genome", a precise quantification of information is difficult to obtain. Some equate information with alleles, some with genes. Some others argue that any genetic variation is information, while others again treat the entire genome of an organism as information. Others still argue that it is Fisher information that increases in natural selection (Frank, 2009). Within the last 15 years a consensus has emerged that the information content of biomolecular sequences (be they protein-coding, RNA coding, or non-coding, such as regulatory motifs) can be measured (or at least estimated) using techniques of Shannon information theory (Adami, 1998; Adami and Cerf, 2000; Adami et al., 2000; Schneider, 2000; Szostak, 2003; Taft and Mattick, 2003; Carothers et al., 2004; Huang et al., 2004; Irizarry et al., 2005; Hazen et al., 2007; Lee et al., 2007; Carbone and Engelen, 2009; Strelioff et al., 2010; Adami, 2012). This information content has an intuitive meaning: it is the amount of information (in bits or any other suitable measure) that the holder has *about* the environment within which it thrives (we will give a mathematical definition below). As an organism uncovers (by the process of evolution) new ways to exploit and survive in that environment, its information about the environment increases. However, changes in the environment can lead to an abrupt (and sometimes catastrophic) decrease in information. Certainly the introduction of anti-viral or anti-bacterial drugs are information-decreasing to the pathogen, and the evolution of drug resistance can be viewed as the attempt of the pathogen to decipher the new environment and acquire the requisite information to exploit it again (Gupta and Adami, 2014). An alternative discussion of evolution as a process to discover information about the environment has been given by Frank (Frank, 2009, 2012b, 2013), who uses a different definition of information.

Measuring information changes along a (historical) evolutionary line of descent is not currently possible, as we would need to have access to extinct homologues of genetic sequences (but see Adami 2012 for a phylogenetic approach). The evidence that information content increases in evolution therefore comes to us via computer simulations (Schneider, 2000) and instantiations of evolution via digital organisms (Adami et al., 2000). In the latter study, a population of asexual digital organisms adapted in a single niche, and a first-order estimate of the information content of the sequences showed that information was steadily increasing as fitness increases. It is clear from this study and others that information has a fitness value (Bergstrom and Lachmann, 2004; Donaldson-Matasci et al., 2010). But is this information actively selected for? To answer this question, we turn to an equation that promises to settle the issue: the Price equation (Price, 1970, 1972a, 1995; Frank, 1995; Gardner, 2008;

Frank, 2012a).

The Price equation is a statistical description of evolutionary change, and relates the changes in a trait variable over time–the left hand side (LHS)–to the covariance of fitness with the trait in question, as well as the expectation of a mutational change in said trait: both terms of the right hand side (RHS) of the equation (the extended Price equation of Kerr and Godfrey-Smith 2008 has three terms). Because of this identity, the Price equation is often described as a tautology (van Veelen, 2005), but it is so only if the LHS and the RHS are related mathematically. If they are related empirically, that is, each of the changes are determined by observing changes in an adapting population, the Price equation can reveal important aspects of the evolutionary dynamics at play. For example, if fitness covaries with the trait, the first term is expected to be positive. Often times, the second term will compensate the first because most mutations on a (selected) trait are deleterious. Naturally, for traits that should be eliminated in order survive, the signs of both terms are reversed. Note that when a population is at mutation-selection balance (no mean change in the trait value), the two terms of the Price equation should cancel each other exactly. While our main focus in this article is to study whether information–treated as a trait value of an adapting organism–is under selection, to some extent we can also answer whether the Price equation is a useful tool to study adaptation.

Treating information as a trait is certainly unconventional. It is clearly not on the same level as other observable (phenotypic) traits. Information is, however, a trait that is uniquely computable from the sequence (at least an estimate of information content is), and would be no different then using the GC content of a sequence (Hildebrand et al., 2010) or a particular transcription factor binding site sequence as a trait that covaries with fitness, or for example. But using information as a trait is different in another respect: Because *everything* that the organisms "knows" about the environment is expressed in this information, it is the trait that should covary the most with fitness. In other words, no other variable should explain as much of the fitness of the organism as the information content. This suggests a very close relationship between information and fitness (Taylor et al., 2007; Polani, 2009; Rivoire and Leibler, 2011), and we will discuss this issue below.

Clearly, using actual evolving populations (such as, for example, the Long-Term Evolution Experiment or LTEE, see Lenski et al. 1991; Lenski and Travisano 1994; Barrick et al. 2009; Wiser et al. 2013) to test the evolution of information would be preferable to all other methods. Unfortunately, as a precise application of Price's equation (see Materials and Methods) requires keeping track of all lineages generation after generation, this is essentially impossible. Our approach to study whether information is a selectable trait is as follows. We will use the digital life platform Avida (Adami and Brown, 1994; Adami, 1998; Ofria and Wilke, 2004; Adami, 2006) to perform evolution experiments that each focus on a different trait, allowing us unprecedented access to empirical data that is not usually available to the experimenter, such as the trait value for every single organism in the entire population, at every single generation. Experiments with 'avidians' have proven to be sophisticated while remaining tractable (see, e.g., Lenski et al. 1999; Wilke et al. 2001; Wilke and Adami 2002; Lenski et al. 2003; Chow et al. 2004; Adami 2006; Misevic et al. 2006; Ostrowski et al. 2007; Clune et al. 2011; Covert et al. 2013; Goldsby et al. 2014).

We study the information content (calculated from the sequence) as one trait, but for the purpose of comparison we will also use another trait (sequence length) that we expect

to be under weak selection, and that can change in either direction. We will compare these experiments to a control treatment where the trait we observe is constructed in such a manner that it cannot be under any positive or negative selection (a neutral marker). In this manner, we can probe the Price equation's ability to distinguish between traits that are under selection as opposed to those that are not, even while the empirical LHS and the RHS must be equal at least statistically, for each and every one of these cases.

We find that information is a strongly selected trait as evidenced by a wide "gap" between the two terms of the Price equation, while sequence length is weakly selected, and the neutral marker is not under selection. While in general a positive covariance between a trait and fitness does not imply a causal relationship between the two observables (either could cause the other, or they could be both caused by a third variable), it is difficult to escape the conclusion that information *causes* fitness. For morphological traits (and even metabolic ones) it is conceivable that high fitness could lead to changes in the trait (rather than the other way around as is more usual). In the present situation, it is hard to imagine how high fitness could cause increased information, as long as the population is well equilibrated. In turn, it is difficult to imagine that a third variable (such as a different trait) causes both high fitness (a possibility) as well as high information (not possible). Thus, we are left with the conclusion that, as intuition would have dictated to us all along, that the origin of an organism's high fitness in a particular environment is the genetic information it has accumulated about that environment.

Methods

The Price Equation

The Price equation (Price, 1970, 1972a, 1995; Frank, 1995, 2012a) provides a statistical link between the mean change in a trait, and population-based averages of the covariance of trait and fitness, as well as the expectation of genetic trait changes (induced by mutations). This equation has been used to address a vast array of topics ranging from the evolution of group selection, kin selection, and cooperation (Wade, 1985; Queller, 1992; Frank, 1995; Page and Nowak, 2002; Henrich, 2004; Fletcher and Zwick, 2006; Okasha, 2006) to ecology (Fox, 2006) and all the way to the selection of music loops (MacCallum et al., 2012). The equation is usually written as follows:

$$\bar{w}\Delta\bar{z} = \mathrm{cov}(w, z) + E(w, \Delta z) \; . \tag{1}$$

with

$$\mathrm{cov}(w, z) = \frac{1}{N_a}\sum_{i=1}^{n_a} n_i(w_i z_i - \bar{w}\bar{z}) \; , \tag{2}$$

$$E(w, \Delta z) = \frac{1}{N_a}\sum_{i=1}^{n_a} n_i w_i(z_i' - z_i) \; . \tag{3}$$

Here, $w_i = \frac{n_i'}{n_i}$ and z_i are the fitness of group i (typically, a genotype) and the trait, respectively, in the ancestral population, n_i and n_i' are the total number of organisms of that type in two

consecutive populations (ancestors and descendants), and $\Delta z_i = z_i' - z_i$ is the amount of change in character in group i between descendant and ancestral (for example by a mutation). The trait could in principle be anything (eye color, resource affinity, mating behavior, etc.) and could be positively or negatively associated with fitness (or it could be entirely neutral). The average fitness in the ancestral population is given by

$$\bar{w} = \frac{1}{N_a} \sum_{j=1}^{n_a} n_j w_j \ , \tag{4}$$

while $\Delta \bar{z} = \bar{z}' - \bar{z}$ is the average amount of change in the trait in two consecutive generations, given by:

$$\bar{z} = \frac{1}{N_a} \sum_{i=1}^{n_a} n_i z_i \ , \quad \bar{z}' = \frac{1}{N_d} \sum_{i=1}^{n_d} n_i' z_i' \ . \tag{5}$$

In the above equations, n_a refers to the number of ancestral types, and $N_a = \sum_{i=1}^{n_a} n_i$ is the ancestral population size (replace 'a' with 'd' for the descendants).

Recently, a generalized form of this equation was derived (Kerr and Godfrey-Smith, 2008) in which immigration and recombination can be considered as genetic sources (in the equation above, the only source of a genotype was from its ancestral type or a mutation). For simplicity, in this work we use the original Price equation as recombination and immigration are absent for the populations we study.

Avida: a Virtual Environment for Digital Organisms

Avida creates a virtual environment within standard desktop computers in which self-replicating computer programs (called "digital organisms" or "avidians") compete for resources, mutate, and evolve (Adami and Brown, 1994; Adami, 1998; Wilke and Adami, 2002; Ofria and Wilke, 2004; Adami, 2006). Avidians can inhabit structured environments (such as would be encountered by bacteria on a Petri dish), or well-mixed environments exemplified by chemostats. As the computer programs must replicate (make copies of the genomic information) to survive, evolutionary 'runs' are experiments much more than simulations, as the information must be physically copied (line by line) from location to location in computer memory, and evolution proceeds via point mutations (and possibly other types such as insertions and deletions) that occur during this process of copying. A point mutation occurs when (with a user-specified probability) a random instruction is written into the daughter genome in place of the instruction that was read.

To be successful, an avidian's growth rate (related to the number of offspring it generates per unit time, the 'update') must exceed the growth rate of other organisms present in the population, and this growth rate depends on the efficiency with which the organism is able to copy its genome, as well as how well it is able to exploit its environment. Avidians obtain the necessary energy to replicate from performing *computations* (which we view as the digital analogue of an organism's metabolism), which allows it to collect resources that can be converted into energy. These computations are mathematical operations performed on random binary numbers that the environment provides, and only a specific sequence of instructions (or variants thereof) trigger the resource and concomitant energy. An increased energy buys

the organism a "faster" CPU (central processing unit), which in turn increases the growth rate. The resources are available at equal concentrations to every organism, approximating a well-stirred chemostat. In the experiments reported here, the resources are unlimited (meaning that they are not depleted by use), which creates a single-niche environment for organisms without negative frequency-dependent selection. In the present experiments, we reward nine different bit-wise logic operations (such as AND, OR, NOR, etc.) an environment commonly referred to as the "logic-9" environment.

In the treatment using information as a trait (as well as for the neutral trait), we limited our population size to 400, which is considerably smaller than the more "standard" population sizes that are an order of magnitude larger. We used smaller population sizes in order to render the computational problem of tracking each individual's lineage in the population tractable. While smaller population sizes create the opportunity for random genetic drift, they are large enough so that we can observe the standard evolutionary dynamics in asexual populations (Elena et al., 2007).

Time in Avida is measured using two different measures. One 'update' has gone by if every organism in the population has executed a fixed number of instructions on average (here set to 30, the default). Because there are 400 organisms in the population, an update has elapsed if 400×30 instructions have been executed. Because it takes organisms hundreds of instructions to create an offspring (this depends on the sequence length, the efficiency of the copying algorithm, and the number of computational tasks executed) many updates can go by before 400 births have taken place, which marks the passing of one generation. Because the Price equation considers subsequent generations when connecting descendant to ancestral populations, we use generations as the unit in time. A complication arises because generations are overlapping in Avida (as they are in most biological populations) so that the same genotype often straddles one (or even more) generations, while the Price equation assumes that, generation by generation, populations are replaced in their entirety. To account for this, when calculating the average change in a character in two consecutive generations, we change the parent identifier (ID) of those descendant entities that straddle both populations (descendant and ancestral) from their original parent ID to their own ID.

The experiments reported in this paper were performed using version 2.12 of the Avida software which is freely available from avida.devosoft.org. The sequence length of the ancestral organism (called 'organism.heads.100') that we used to seed every reported run is 100 instructions. This simple organism can self-replicate but is not able to accomplish any logical computations. For more details about Avida, including the default set of instructions and the definition of the environment, refer to (Ofria and Wilke, 2004), and to the web resource above. Experiments ran under the Linux operating system on an Intel10 1536 cluster of 192 dual core Intel 2200+ processors at Michigan State University's High Performance Computing Centre.

Fitness and Information Content of Digital Organisms

The fitness of a genotype w in a particular environment is defined by a typical representative's *expected* rate of offspring production. This definition makes it easy to apply the Price equation to a population of Avidians. Fitness can be *estimated* by the ratio of necessary energy for an organism to execute the instructions in its genome and the required number of instructions to produce an offspring (Adami, 1998; Elena et al., 2007), but the realized fitness depends on

the organism's interaction with others and the vagaries of chance. The mandatory energy for reproduction is provided to a digital organism in the form of discrete quanta, called "Single-Instruction Processing" units, or SIPs (Lenski et al., 2003; Adami, 2006; Elena et al., 2007). The rate at which an organism acquires energy depends on its genome length and a coefficient which is related to the interactions between the organism's computational metabolism and its environment. In other words, an organism gets more SIPs if it accomplishes specific computational functions in an environment that rewards those functions. As a consequence, it is easy to compare the relative rates of reproduction of any two given organisms. If an organism has w twice as large as another organism, then it is expected that the former organism produces offspring at twice the rate of the latter. Moreover, organisms that cannot produce any offspring have zero fitness. The absolute fitness of the ancestral organism ('organism.heads.100') is 0.2494, which we use to normalize the fitness of the evolved populations relative to this ancestor. Note that this fitness is calculated by allowing the organism to grow in a single virtual dish, and so is an estimator of long-term success. During an experiment, fitness is not calculated for any organisms: either they survive and leave offspring or they do not. When calculating the information content of single sequences as described below, we have to estimate the fitness of all possible one-mutants of that sequence, which is again obtained via virtual growth-assays in isolation.

In general, information is defined by Shannon as the shared entropy between two ensembles (Shannon, 1948; Cover and Thomas, 2006) and therefore cannot be determined for a single sequence. To circumvent this, we use a form of information that is conditional on a fixed environment (the environment that the organism actually encounters). For an organism with a fixed sequence length ℓ, we can calculate the maximal entropy (uncertainty) for a hypothetical infinite population: this would be $H(\max) = \ell$ (if we agree to take logarithms to the base D, where D is the number of possible instructions, that is, the alphabet size) because the likelihood to find any particular sequence in such an ensemble is $p = 1/D^\ell$. In the present experiments, $D = 26$.

This maximal entropy only describes *unselected* (that is, random) sequences. Those that accomplish a function (and thus replicate) have a very different entropy $H_{\rm actual}$, which depends on the relative concentration of each of the existing types in the population. This entropy depends on the environment (thus it is a conditional entropy). If we associate the genetic sequence of a particular organism with the states of a random variable A, and describe the environment as the specific state $E = e$ of an environment random variable, we can write the information content of organism A as (Adami, 1998; Adami and Cerf, 2000)

$$I(A) = H(A) - H(A|E = e) \; , \tag{6}$$

where we have followed the convention that unconditional entropies are maximal, i.e., $H(A) = H_{\rm max} = \ell$) and

$$H(A|E = e) = -\sum_a p(a|e) \log_D p(a|e) \; , \tag{7}$$

where $p(a|e)$ is the probability to find the self-replicating type $A = a$ given that the environment is in state $E = e$. The actual entropy $H(A|e)$ in Eq. (7) is in principle defined via a sum over an infinite ensemble of sequences (all sequences a that, given e, are functional).

We will simulate such infinite ensembles artificially via an exhaustive mutational process (described below), because sampling the frequencies of organisms in an actual (small) evolving population cannot approach the true entropy (7) (Huang et al., 2004).

Generally speaking, the entropy of a polymer can be written in terms of the sum of the entropy of the monomers (at each site), the pairwise entropies (shared entropy between sites), shared entropy between triples of sites (and so on), using a formula due to Fano (1961). Writing the random variable A in terms of the monomeric random variables $A^{(i)}$

$$A = A^{(1)} \times A^{(2)} \times ... \times A^{(\ell)} \tag{8}$$

the entropy $H(A|e)$ can be written as (Fano, 1961)

$$H(A|e) = \sum_{i=1}^{\ell} H(A^{(i)}|e) - \sum_{i<j}^{\ell} H(A^{(i)} : A^{(j)}|e) + \sum_{i<j<k}^{\ell} H(A^{(i)} : A^{(j)} : A^{(k)}|e) - \cdots \tag{9}$$

In (9), the sum goes over alternating signs of correlation entropies, culminating with a term $(-1)^{\ell-1} H(A^{(1)} : A^{(2)} : \cdots : A^{(\ell)}|e)$. Here, we neglect all higher-order correlation terms and assume that the entropy $H(A|e)$ can be approximated by the sum over the per-site entropies $\sum_{i=1}^{\ell} H(A^{(i)}|e)$. In that case, Eq. (6) can be written simply as

$$I(A) = \ell - \sum_{i=1}^{\ell} H(A^{(i)}|e) \; . \tag{10}$$

The entropies $H(A^{(i)}|e)$, in turn, can be obtained by creating all single mutants of organism A, and measuring their fitness in isolation. We count all single mutations at site i that are either neutral (or in rare cases beneficial) to obtain ν_i and estimate the entropy of site i as $\log_D \nu_i$, as in (Adami et al., 2000). By taking logarithms to the base D, we normalize the entropy in such a way that a single monomer (a single instruction) has maximal entropy 1 'mer' (Adami, 2004). Naturally, the maximal entropy of an ℓ-mer then is ℓ mers, which is also the maximal information (for $D = 26$, one mer equals approximately 4.7 bits).

Neglecting the sub-leading terms $\sum_{i<j}^{\ell} H(A^{(i)} : A^{(j)}|e)$ can have significant repercussions (see, e.g., Gupta and Adami 2014), but we do not expect them to affect the conclusions we reach about selection even though they will affect the quantitative estimate of the information content.

Results

We test whether three different traits (information, sequence length, and a neutral marker) are evolutionarily selected, by measuring the mean rate of change of the trait (the LHS) of the Price equation, and in particular the two terms on the RHS of the Price equation.

Information as a Selectable Trait

We set up an experiment to record fitness and information content (as the trait) for every organism in the population at every generation as described in Methods, for 14,000 generations. The sequence length is fixed to remain at 100 instructions, while the point mutation rate is set to $\mu = 0.015$ so that on average each genome receives 1.5 mutations per generation. Insertion and deletion mutations are turned off. We chose the relatively high mutation rate so that significant evolutionary changes can be obtained with a comparably short experiment of 14,000 generations (still requiring 5.6×10^6 trait evaluations, each of which requires 2,600 fitness evaluations: one for each substitution at each of the 100 sites, for a total of over 10 billion fitness evaluations per run).

We show in Figure 1 the mean fitness and average information content over 14,000 generations. Fitness shows the usual stair-step increase (note that we plot the logarithm of fitness). The 100 instruction-long ancestral sequence consists of a self-replicating 'gene' that spans 15 instructions, and 85 inert (meaning, non-functional) `nop-C` instructions. Because the fitness of the organism is unlikely to change if the inert instruction is replaced by any of the other D instructions, each contributes $\log \nu = \log D = 1$ to the conditional entropy $H(A|e)$, which subtracts from the first term in (10). Thus, the information content begins at about 15 (it is slightly lower because the replication gene has some redundancy in it). Information content increases monotonically at first as more and more of the 85 "empty" instructions are filled with information, but later we can see periods where the measured information decreases.

We attribute these decreases to our use of the rough estimate (10) as a proxy for the full information content that uses Eq. (9). While this approximation introduces noise, we are forced to do so because even including the first correction term $\sum_{i<j}^{\ell} H(A^{(i)} : A^{(j)}|e)$ would make the computation almost intractable. Fluctuations in average information content are likely due to the high mutation rate.

Using the average *change* in information in subsequent generations $\Delta\bar{z}$ and a calculation of the two terms on the RHS of the Price equation (1) directly from the population, we first test whether the Price equation is obeyed. Fig. 2 confirms that this is the case: we plot the LHS vs. the RHS, and note that they cluster tightly around the identity line. We also note that if the Price equation were a mere tautology, we should obtain a perfectly straight line. Instead, we observe that the data cluster slightly below the identity line, which is likely an artifact caused by the small population size (the bias is absent in the second treatment where the population size is much larger). In the inset of Fig. 2 we show that the distribution of the LHS is slightly skewed to the right, implying that there are more transitions towards higher information than there are towards lower information. This observation simply corroborates that, on average, information is increasing.

We now turn to analyzing the LHS and the RHS of Eq. (1) separately. The first term, which also appears in the Price equation without mutations, quantifies the covariance of fitness with the trait. Indeed, in the absence of mutations, it is this term that drives the mean change of trait (the LHS) positive or negative (depending on whether the covariance is positive or negative). It thus reflects the force of purifying selection: it vanishes once all variance in the population has disappeared. Note that if fitness was used as a trait, the covariance would simply be the variance of fitness (the covariance of fitness with itself), and the Price equation would become Fisher's Fundamental Theorem as has been noted repeatedly, not the least by

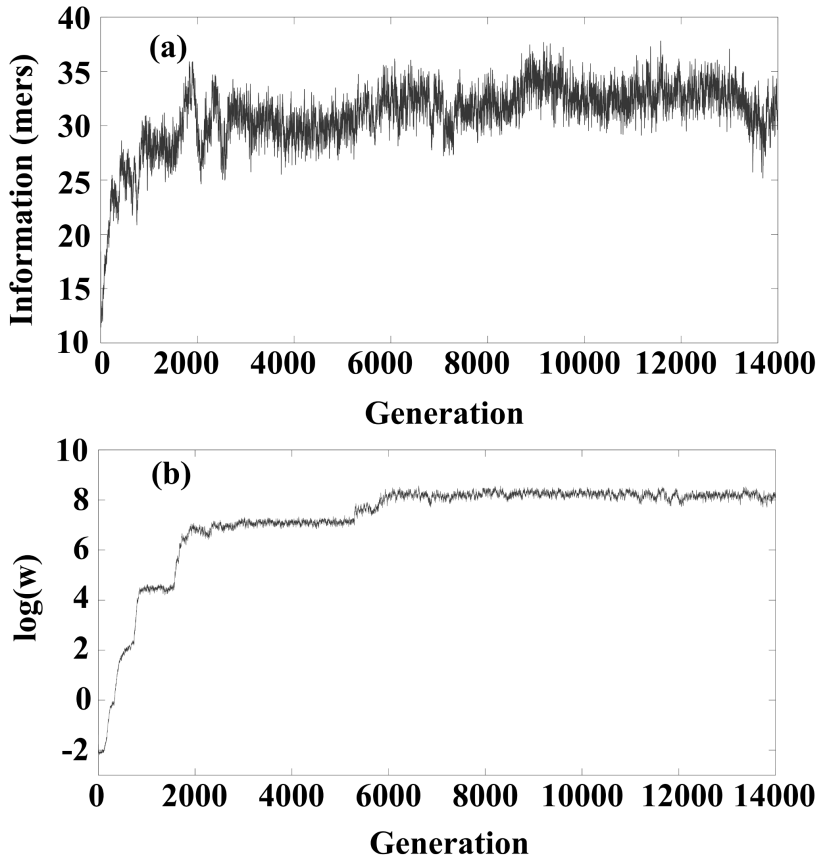

Figure 1: Average fitness and information content in a population of avidians through evolution. (a): average population information content. (b): Average absolute fitness.

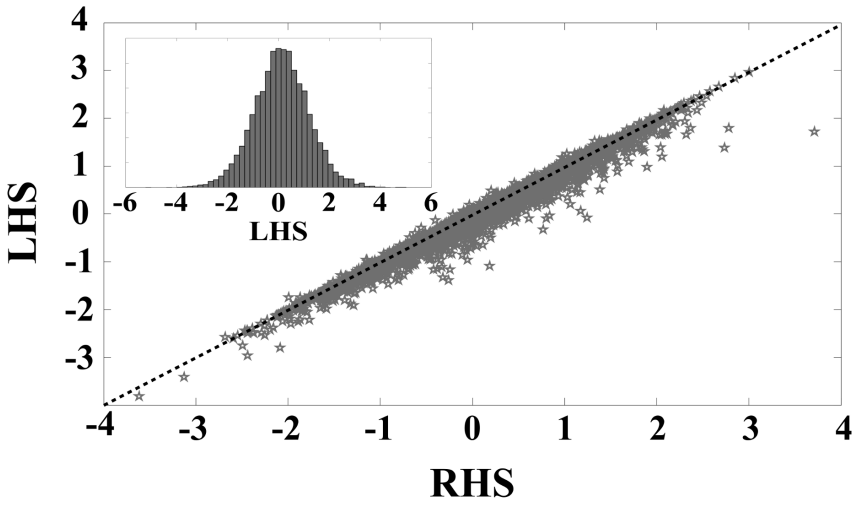

Figure 2: Plot of the LHS vs. the RHS of the Price equation (1) with information as trait. The dashed line is the identity line, and the inset show the distribution of the LHS of the Price equation), which is approximately normal with a mean shifted to the right of zero.

Price himself (Price, 1972b).

The second term of the RHS takes mutations into account explicitly, as without mutations $\Delta z = 0$ (the trait could not change from generation to generation, only the frequency of traits could). It is this term that drives adaptation: single rare mutations can give rise to large positive Δz. However, in most cases, Δz will be negative, as mutations disrupt the trait. Thus, the distribution of Δz reflects the distribution of beneficial, neutral, and deleterious mutations. And because most mutations are detrimental, the second term of the RHS is mostly negative even though the trajectory of increasing $\bar{z}$ is driven by the rare beneficial mutations.

In Fig. 3a we show the size of first and second term of the RHS of the Price equation separately, and see that indeed the first term is positive on average, and the second term is negative on average. It is selection that drives these terms apart (as we will also show explicitly in the second treatment and the control): the first term shows strong covariance of information with fitness, and the second term shows that this positive association is balanced by mutations that tend to destroy the trait. In Fig. 3b we plot the LHS against the first and second term separately, and show that neither lies on the identity line by themselves. The clear separation between the first term and the second term–which we term the "Price equation gap"–is an indication of the strength of selection operating on the trait. We would like to stress that using the first term alone to assess strength of selection on a trait is not sufficient, as *any* trait is likely to change as the population approaches a pure (meaning homogeneous) state. The Price equation gap, on the other hand, is a strong indicator as it reflects the forces of selection in mutation-selection balance.

Weak Selection on Sequence Length

In order to be able to appreciate the nature of selection on information, we also study selection on a more "classical" trait (the sequence length) using the Price equation. In the first treatment, sequence length was kept constant so as not to confound selection on length with selection on information. For this control treatment, we allow sequence length to change using insert and delete mutations at a rate of 0.01 and 0.001 per genome per generation respectively, along with the copy-mutations at a rate of $\mu = 0.00175$ per instruction copied, and a constant population size of 1,400 individuals. Note that the ratio of insert to delete mutations is biased towards an increase in sequence length if the sequence was drifting neutrally.

We seed the experiment with the same ancestor as the previous experiment, but expect sequence length changes to occur that are adaptive. While extra instructions are not directly detrimental (or beneficial) to an individual's fitness because each organism obtains SIPs in proportion to sequence length (Ofria and Wilke, 2004), there is a weak selection pressure to reduce sequence length so as to reduce mutational load (Lenski et al., 1999). There is also a pressure to increase sequence length to accommodate new beneficial genes, but only if the current sequence length is insufficient to accommodate the new information. Thus, we expect to see selection in both directions, and indeed we witness first a trend towards reduced sequence length (as this allows the population to reduce the mutational load) and later intermittent trends towards higher sequence length (see Fig. 4a). Fitness is steadily increasing during this period of changing sequence length, as evidenced by Fig. 4b.

We record all individual genomes (along with the trait under consideration, the sequence

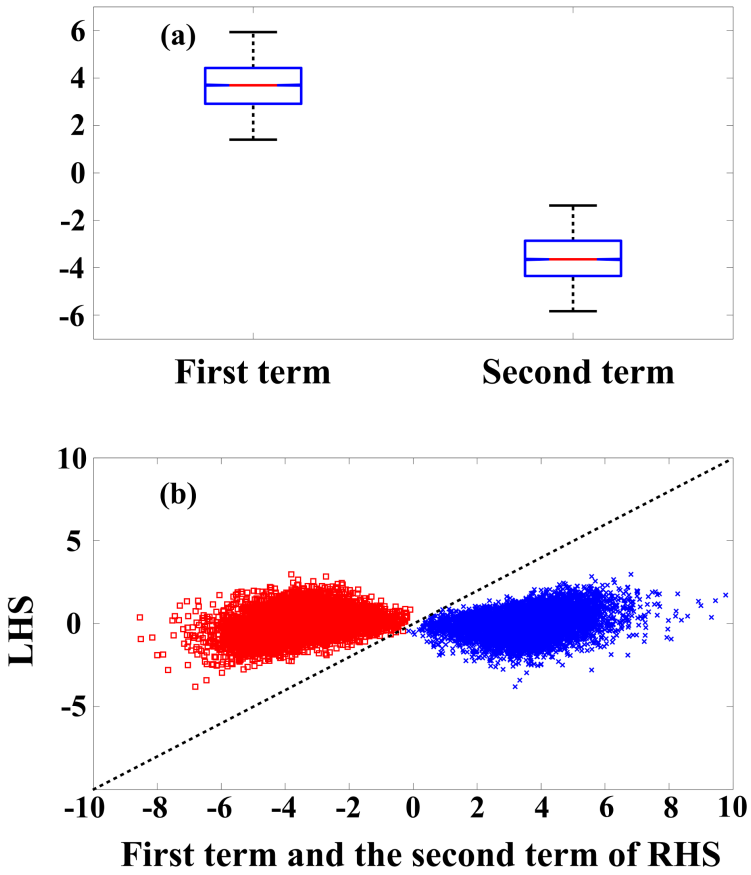

Figure 3: Opposing contributions of first and second term of RHS of Price equation. (a): Size of first and second term, averaged over the 14,000 generations of a typical experiment. The box represents the first and third quartile while the red line represent the median, and the whiskers encompass 99% of the data. Notches (blue) indicate 95 percent confidence interval of the median. (b) LHS of Price equation vs. first (blue dots) and second (red dots) term of the RHS separately (we show a random sample of 7,000 out of 14,000 points), showing that the counterbalancing of two terms is necessary to satisfy the Price equation.

length), and their fitness for 14,000 generations, and calculate mean changes in trait every generation (for the LHS), as well as the covariance and the expectation of trait change (for the two terms of the RHS). We first verify that the Price equation is obeyed by plotting LHS vs. RHS (see Fig. 5). Note that the distribution of length changes (shown in the inset of Fig. 5) is slightly biased towards decreases.

Next, we study the Price equation gap by plotting in Fig. 6a the size of the first term and the second term of the RHS individually as before. Both the first and second term medians are significantly different from zero (the notches in the box plot are 95% confidence intervals) and from each other, but the effect is small. The overall Price equation gap (averaged over the entire evolutionary history) is small. Plotting the LHS vs. the first and second terms of the RHS independently (as in Fig. 3b) shows that the two terms are not well separated, and instead overlap significantly. The mutation term (red dots in Fig. 6b) seems to skew towards

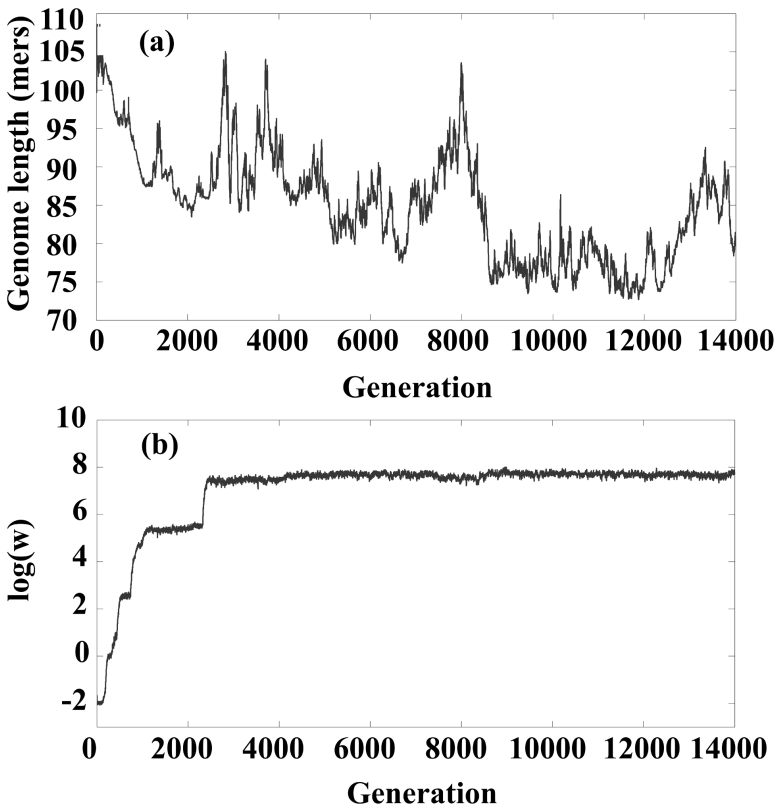

Figure 4: Population-averaged sequence length and fitness. (a): Mean sequence length as a function of time (in generations). (b): Mean absolute fitness as a function of time (in generations).

positive changes, which could be attributed to the biased insert mutation rate rather than to a selective pressure favoring increases in the trait. We conclude that sequence length is under selection, but only weakly so (the first term is weakly negative, and the second term is weakly positive) when averaged over the entire trajectory. The overall signal of selection strength can be weakened when selection acts in opposing directions at different times, and we should expect to see stronger signals if we would focus on those parts of the trajectory where sequence length is exclusively under positive or negative selection.

Price Equation for a Neutral Marker

As a control for the two previous treatments we test the Price equation on a marker that is by construction completely neutral. In this experiment, we study as a trait the difference in frequency between two specific neutral "alleles", the instructions `nop-A` and `nop-B`. These "no operation" instructions have a role in the function of a typical avidian genome as they can specify the register that is used in internal calculations of the CPU, or else they can mark areas of code for recognition (pattern-based addressing). However, the frequency of this instruction is not at all under selection, and adding or deleting this instruction in code that is not specifically involved in computations is entirely neutral. In particular, the *difference* in frequency between those two instructions is expected to be unselected even if the absolute

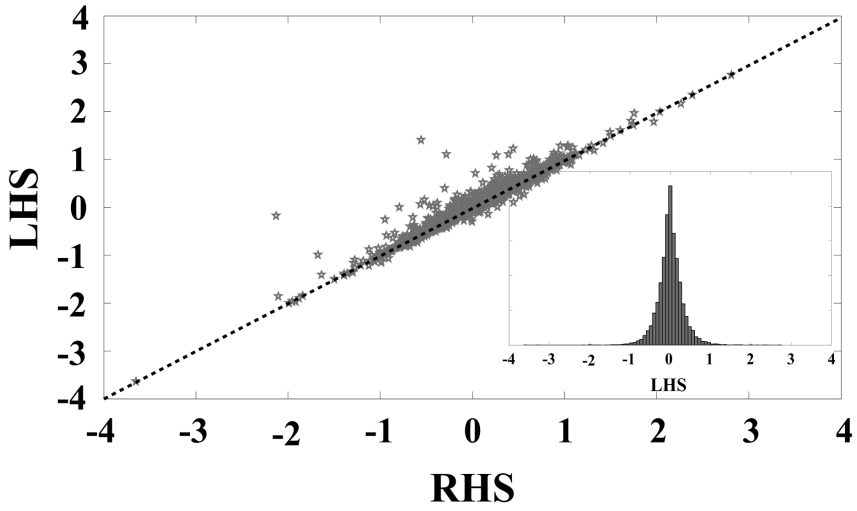

Figure 5: LHS vs. RHS of the Price equation for sequence length as trait. Each dot represents a (LHS,RHS) pair along the evolutionary trajectory (we show 7,000 randomly selected points out of 14,000). The dashed black line represents the identity line. The inset shows the distribution of trait changes (the LHS).

value of either was. Note that we did not use the third no-instruction operation `nop-C` as a neutral marker as the ancestral genome is filled with 85 of those instructions, so reducing this number will be under strong selection initially. Because we collected the data for the neutral trait at the same time as we collected the information trait, we can use the same experiment.

As expected, the trait does not change significantly during evolution, as a plot of the LHS vs. the RHS of the Price equation for this trait reveals (see Fig. 7). We can also study the Price equation gap, and find it to be non-existent (see Fig. 8a), but not just on average, but for the entire experiment. There is little variation in the LHS or RHS, and no difference between the two terms of the RHS, and with zero. Plotting the two terms of the RHS against the LHS (Fig. 8b) reveals nothing. Indeed, the second term of the RHS (the expectation of trait change) is almost never negative, indicating that mutations on this trait are never deleterious. All observed changes in this trait are consistent with neutral drift.

Discussion

"Information" is a modern buzzword. Some go so far as to claim that "everything is information" (Wheeler, 1990), but certainly within molecular biology the idea that information is stored in our genes has strong support since the discovery of DNA. But it certainly has not always been clear how we should measure the amount of information coded in genes, and in particular whether and how selection could act on it. Clearly, information is not the same thing as "the set of all nucleotides" within the chromosomes of an organism: information must be a "difference that makes a difference" (Bateson, 1972) (or at least, that *can* make a difference). Shannon's theory of information (Shannon, 1948; Cover and Thomas, 2006) provides precisely the framework that is necessary to tell the difference between sequence and

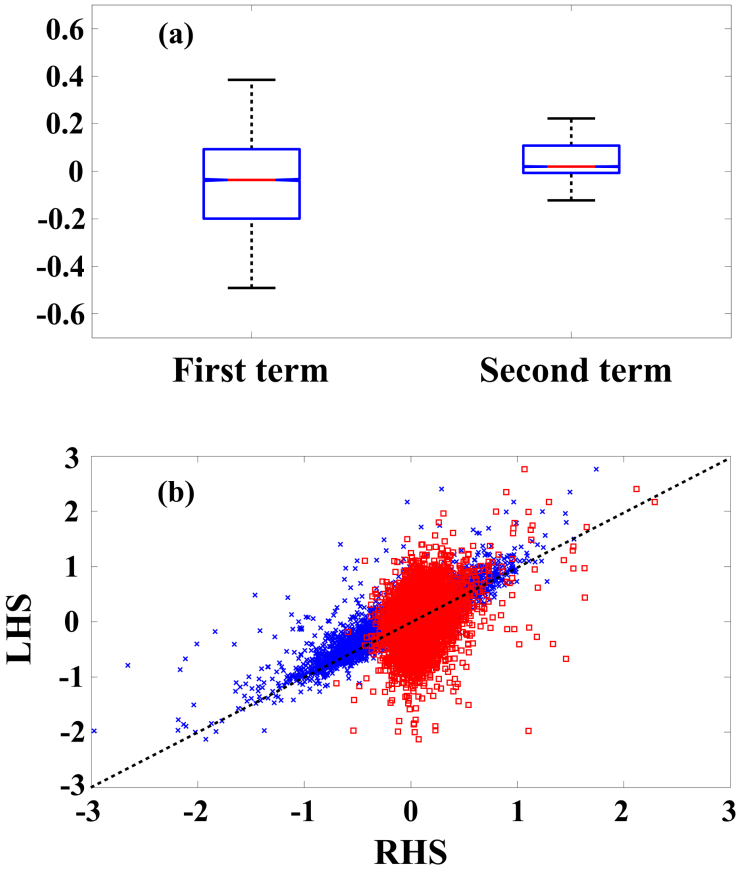

Figure 6: Contributions to the RHS of the Price equation for sequence length trait. (a): A comparison of the size of the first and second term shows that they are significantly different from zero and from each other. The box again represents the first and third quartile while the red line represent the median, and the whiskers encompass 99% of the data. Notches (blue) indicate 95 percent confidence interval of the median. Note the small range of the trait. (b): Scatter plot of LHS vs. first (blue points) and second term (red points) of Price equation individually shows that the distributions overlap (random sample of 7,000 out of 14,000 points shown).

information: information is that which allows the bearer to make predictions with accuracy better than chance. Organisms constantly make predictions about their environment (even if this is not always immediately apparent). The metabolic genes, for example, make predictions about the nature and form of resources that the organism is likely to encounter, and developmental genes "predict" how to make the organism using only the genetic code and the environment it finds itself in. A gene in males that encodes a trait that leads to sexual selection predicts that there are females around that are susceptible to this trait. In a sense then, it may be that all traits that have a selective value are informative. But how can we measure this information? We have previously proposed that information can be measured in populations that adapt to an environment by measuring the frequency of substitutions at individual nucleotide positions: those that do not encode information are not under selection (and therefore should be randomly drifting), while those that encode information should

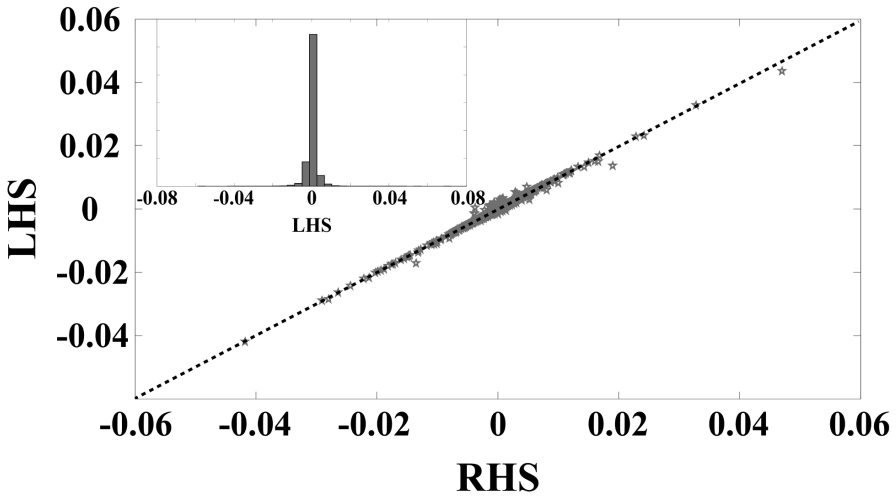

Figure 7: LHS vs. RHS of the Price equation for a neutral trait. The Price equation is strictly obeyed, but the changes in trait value (LHS) as well as the covariance and expectation of trait change are much smaller than in the previous treatments (note the scale on the axes). The distribution of trait changes (see inset) is symmetric and strongly peaked at zero.

be fixed as any mutations would have severe fitness consequences (Adami, 1998; Adami and Cerf, 2000; Adami et al., 2000; Adami, 2002, 2004). Such a measure was previously applied to DNA binding motifs (Schneider et al., 1986) (where the problem of possible higher-order interactions between sites is much reduced) and is used ubiquitously to display sequence motifs (Schneider and Stephens, 1990). While this method is most easily used to study how information evolves in computational simulations of evolution (Schneider, 2000; Hintze and Adami, 2008) or digital evolution (Adami et al., 2000; Ofria et al., 2008), the usefulness of this measure in characterizing the evolving functionality of biomolecules (as opposed to binding sites or digital molecules) has also been shown repeatedly (Carothers et al., 2004; Materon et al., 2004; Carothers et al., 2006b,a; Hazen et al., 2007; Gupta and Adami, 2014).

It is clear that natural selection (acting on fitness) is responsible for creating the allele frequency differences that are ultimately responsible for the information, and it could be argued that, as a consequence, information and fitness are one and the same thing. They are not, as fitness is measured in offspring per unit time and information is dimensionless and measured in bits (or any other convenient unit). There is no doubt, however, that there is a deep link between the two, and recent work has progressed in unraveling that link (Rivoire and Leibler, 2011). Here, we decided to test explicitly if information–defined as a measurable (and thus phenotypic) trait–is under selection. In principle, this information is independent of the actual sequence that encodes the information (the same information could be encoded by very different sequences) while in practice we find that the same information is encoded by similar sequences, most likely because inheritance enforces this. Information cannot be broken down to "letters": the information content of a single sequence is not defined, unless that sequence is either compared to a specific ensemble (Schneider, 1997), or if an ensemble of sequences is created from that sequence, as we have done here. Of course, each individual trait that contributes to fitness has an underlying genetic encoding, and each such trait has

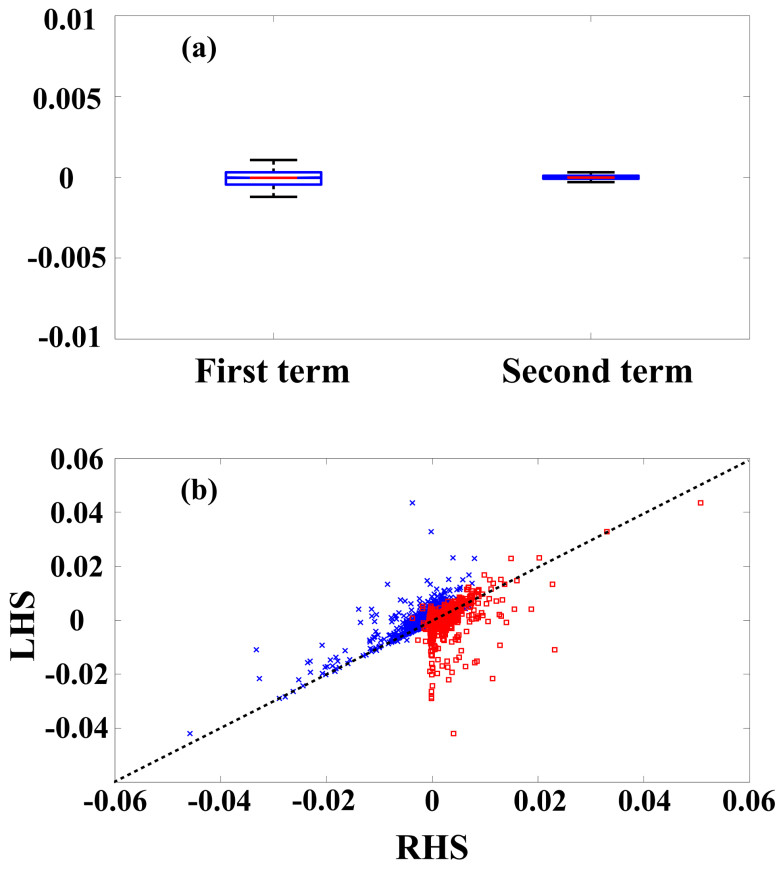

Figure 8: Contributions to the RHS of the Price equation for a neutral marker trait. (a): A comparison of the size of the first and second term shows that they are not significantly different from zero and from each other. Boxplot as in Fig. 6b. (b): Scatter plot of LHS against first term of RHS (blue) and second term of RHS (red) of the Price equation for the neutral trait.

an informational value. The avidians we have used for the experiments we reported here also have individual traits (their capacity to carry out any of the nine different logical functions, or aspects of their gene sequence for self-replication), and the information necessary to carry out each task can be measured individually (Ofria et al., 2008). But the total information content is a complicated function of the information encoding each individual trait, so it is possible that the total information content is the single best predictor of an organism's success. Indeed, this may ultimately turn out to be the enduring advantage of using an information-based description: if information is the best proxy for fitness, then measuring information is sufficient to infer fitness (which by itself is of course notoriously difficult to measure or estimate, see, e.g., Orr 2009).

By using information as trait value, we can study the degree to which that trait is under selection using the Price equation. We acknowledge that the value of the Price equation–as wells as its meaning–have been discussed at length in the literature (see, e.g, the review by Frank 2012a). We corroborate here that the Price equation is indeed an identity that "partitions evolutionary change into two components" (Frank, 2012a), and that what makes the equation useful is that the two components play very different roles, and are sensitive

to the dynamics. For example, during evolutionary stasis the two terms must cancel each other exactly, whereas during periods of intense positive or negative selection, they take on values significantly different from their stasis values. This is true in the experiments we have performed here too: we can look at the first and second terms of the Price equation averaged over parts of the evolutionary trajectory (a selective sweep or a period of stasis) and notice dramatic changes in the mean values of the first and second term over that period (data not shown). We also stress that the monotonic increase in information can only hold in a fixed environment. An abrupt change in the environment is almost always going to lead to an abrupt *drop* in information, as what used to be information (predicting the previous environment's features) may after the change have turned to entropy. We expect after the change to witness negative covariance of the "old" information with fitness, and a positive second term, as mutations create new information to replace the old.

While information content (as a phenotypic trait) has a number of appealing characteristics, we hasten to add that there are numerous caveats associated with it, which (at least today) give fitness measurements the edge. For estimating information content, we cannot rely on the substitution patterns obtained from small populations, as the shared descent of molecules creates a spurious signature of information (Huang et al., 2004). To remedy this, we must generate artificial ensembles by creating all point-wise single mutants, double mutants, and so on, of a sequence. Neglecting the high-order correlations between sites is the only feasible computational shortcut, but this can lead to important biases (mostly under-estimating the information content, see Gupta and Adami 2014). However, we believe that thinking in terms of information rather than in terms of fitness can help us navigate the sometimes bewildering jungle of evolutionary genetics, and put evolutionary theory on an even more solid quantitative footing.

Acknowledgements

We thank Charles Ofria and Anurag Pakanati for extensive discussions on information estimation, and Thomas LaBar for insightful comments on the manuscript. This work was supported in part by the National Science Foundations BEACON Center for the Study of Evolution in Action, under Contract No. DBI-0939454. We wish to acknowledge the support of the Michigan State University High Performance Computing Center and the Institute for Cyber Enabled Research.

References

Adami, C. 1998: Introduction to Artificial Life. Springer-Verlag New York, Inc., New York, NY, USA.

———. 2002: What is complexity? BioEssays 24(12):1085–1094.

———. 2004: Information theory in molecular biology. Phys. Life Reviews 1:3–22.

———. 2006: Digital genetics: unravelling the genetic basis of evolution. Nature Reviews Genetics 7(2):109–118.

———. 2012: The use of information theory in evolutionary biology. Annals NY Acad. Sci. 1256:49–65.

Adami, C. and Brown, C. 1994: Evolutionary learning in the 2D Artificial Life system Avida. Proceedings of the 4th International Conference on the Synthesis and Simulation of Living Systems (Artificial Life 4), (edited by R. Brooks and P. Maes), pages 377–381. MIT Press.

Adami, C. and Cerf, N. 2000: Physical complexity of symbolic sequences. Physica D 137(1):62–69.

Adami, C.; Ofria, C.; and Collier, T. C. 2000: Evolution of biological complexity. Proceedings of the National Academy of Sciences of the U.S.A. 97(9):4463–4468.

Barrick, J. E.; Yu, D. S.; Yoon, S. H.; Jeong, H.; Oh, T. K.; Schneider, D.; Lenski, R. E.; and Kim, J. F. 2009: Genome evolution and adaptation in a long-term experiment with Escherichia coli. Nature 461(7268):1243–7.

Bateson, G. 1972: Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology. Chandler Publishing Company.

Bergstrom, C. and Lachmann, M. 2004: Shannon information and biological fitness. Proceedings IEEE Workshop on Information Theory, pages 50–54. IEEE.

Carbone, A. and Engelen, S. 2009: Information content of sets of biological sequences revisited. Algorithmic Bioprocesses (Natural Computing Series), (edited by A. Condon; D. Harel; J. N. Kok; A. E Salomaa; and E. Winfree), pages 31–42. Springer Verlag, Berlin, Heidelberg.

Carothers, J. M.; Davis, J. H.; Chou, J. J.; and Szostak, J. W. 2006a: Solution structure of an informationally complex high-affinity RNA aptamer to GTP. RNA 12(4):567–79.

Carothers, J. M.; Oestreich, S. C.; Davis, J. H.; and Szostak, J. W. 2004: Informational complexity and functional activity of RNA structures. J. American Chem. Society 126:5130–5137.

Carothers, J. M.; Oestreich, S. C.; and Szostak, J. W. 2006b: Aptamers selected for higher-affinity binding are not more specific for the target ligand. J Am Chem Soc 128(24):7929–37.

Chow, S.; Wilke, C.; Ofria, C.; Lenski, R.; and Adami, C. 2004: Adaptive radiation from resource competition in digital organisms. Science 305(5680):84–86.

Clune, J.; Goldsby, H. J.; Ofria, C.; and Pennock, R. T. 2011: Selective pressures for accurate altruism targeting: evidence from digital evolution for difficult-to-test aspects of inclusive fitness theory. Proc Biol Sci 278(1706):666–74.

Cover, T. and Thomas, J. 2006: Elements of Information Theory. Wiley-interscience, New York, NY.

Covert, A. W., 3rd; Lenski, R. E.; Wilke, C. O.; and Ofria, C. 2013: Experiments on the role of deleterious mutations as stepping stones in adaptive evolution. Proc Natl Acad Sci U S A 110(34):E3171–8.

Donaldson-Matasci, M.; Bergstrom, C.; and Lachmann, M. 2010: The fitness value of information. Oikos 119(2):219–230.

Elena, S.; Wilke, C.; Ofria, C.; and Lenski, R. 2007: Effects of population size and mutation rate on the evolution of mutational robustness. Evolution 61(3):666–674.

Fano, R. M. 1961: Transmission of Information: A Statistical Theory of Communication. MIT Press and John Wiley, New York and London.

Fletcher, J. A. and Zwick, M. 2006: Unifying the theories of inclusive fitness and reciprocal altruism. Am Nat 168(2):252–62.

Fox, J. W. 2006: Using the Price equation to partition the effects of biodiversity loss on ecosystem function. Ecology 87(11):2687–96.

Frank, S. 1995: George Price's contributions to evolutionary genetics. Journal of Theoretical Biology 175(3):373–388.

Frank, S. A. 2009: Natural selection maximizes Fisher information. J Evol Biol 22(2):231–44.

———. 2012a: Natural selection. IV. The Price equation. J Evol Biol 25(6):1002–19.

———. 2012b: Natural selection. V. How to read the fundamental equations of evolutionary change in terms of information theory. J Evol Biol 25(12):2377–96.

———. 2013: Natural selection. VI. Partitioning the information in fitness and characters by path analysis. J Evol Biol 26(3):457–71.

Gardner, A. 2008: The Price equation. Current Biology 18(5):R198–R202.

Goldsby, H. J.; Knoester, D. B.; Ofria, C.; and Kerr, B. 2014: The evolutionary origin of somatic cells under the dirty work hypothesis. PLoS Biol 12(5):e1001858.

Gupta, A. and Adami, C. 2014: Changes in epistatic interactions in the long-term evolution of HIV-1 protease. ArXiv.org preprint arXiv:1408.2761.

Hazen, R. M.; Griffin, P. L.; Carothers, J. M.; and Szostak, J. W. 2007: Functional information and the emergence of biocomplexity. Proc. Natl. Acad. Sci. USA 104:8574–8581.

Henrich, J. 2004: Cultural group selection, coevolutionary processes and large- scale cooperation. J. Econ. Behav. Organiz. 53:3–35.

Hildebrand, F.; Meyer, A.; and Eyre-Walker, A. 2010: Evidence of selection upon genomic GC-content in bacteria. PLoS Genet 6(9):e1001107.

Hintze, A. and Adami, C. 2008: Evolution of complex modular biological networks. PLoS Computational Biology 4(1):e23.

Huang, W.; Ofria, C.; and Torng, E. 2004: Measuring biological complexity in digital organisms. Proceedings of Artificial Life IX, (edited by J. Pollack; M. A. Bedau; P. Husbands; T. Ikegami; and R. Watson), pages 315–321. MIT Press, Cambridge, MA.

Irizarry, K. J. L.; Merriman, B.; Bahamonde, M. E.; Wong, M.-L.; and Licinio, J. 2005: The evolution of signaling complexity suggests a mechanism for reducing the genomic search space in human association studies. Mol Psychiatry 10(1):14–26.

Kerr, B. and Godfrey-Smith, P. 2008: Generalization of the Price equation for evolutionary change. Evolution 63(9999):531–536.

Lee, C.-T.; Risom, T.; and Strauss, W. M. 2007: Evolutionary conservation of microRNA regulatory circuits: an examination of microRNA gene complexity and conserved microRNA-target interactions through metazoan phylogeny. DNA Cell Biol 26(4):209–18.

Lenski, R. E.; Ofria, C.; Collier, T. C.; and Adami, C. 1999: Genome complexity, robustness and genetic interactions in digital organisms. Nature 400:661–664.

Lenski, R. E.; Ofria, C.; Pennock, R. T.; and Adami, C. 2003: The evolutionary origin of complex features. Nature 423(6936):139–144.

Lenski, R. E.; Rose, M. R.; Simpson, S. C.; and Tadler, S. C. 1991: Long-term experimental evolution in Escherichia coli. I. Adaptation and divergence during 2,000 generations. Am Nat 138:1315–1341.

Lenski, R. E. and Travisano, M. 1994: Dynamics of adaptation and diversification: a 10,000-generation experiment with bacterial populations. Proc Natl Acad Sci U S A 91(15):6808–14.

MacCallum, R. M.; Mauch, M.; Burt, A.; and Leroi, A. M. 2012: Evolution of music by public choice. Proc Natl Acad Sci U S A 109(30):12081–6.

Materon, I. C.; Beharry, Z.; Huang, W.; Perez, C.; and Palzkill, T. 2004: Analysis of the context dependent sequence requirements of active site residues in the metallo-beta-lactamase IMP-1. J Mol Biol 344(3):653–63.

Misevic, D.; Ofria, C.; and Lenski, R. 2006: Sexual reproduction reshapes the genetic architecture of digital organisms. Proceedings of the Royal Society B: Biological Sciences 273(1585):457–464.

Ofria, C.; Huang, W.; and Torng, E. 2008: On the gradual evolution of complexity and the sudden emergence of complex features. Artif Life 14(3):255–63.

Ofria, C. and Wilke, C. O. 2004: Avida: A Software Platform for Research in Computational Evolutionary Biology. Artificial Life 10(2):191–229.

Okasha, S. 2006: Evolution and the Levels of Selection. Oxford Univ Press, Oxford, UK.

Orr, H. A. 2009: Fitness and its role in evolutionary genetics. Nat Rev Genet 10(8):531–9.

Ostrowski, E.; Ofria, C.; and Lenski, R. 2007: Ecological specialization and adaptive decay in digital organisms. The American Naturalist 169(1):E1–E20.

Page, K. M. and Nowak, M. A. 2002: Unifying evolutionary dynamics. J Theor Biol 219(1):93–8.

Polani, D. 2009: Information: currency of life? HFSP J 3(5):307–316.

Price, G. R. 1970: Selection and covariance. Nature 227:520–521.

———. 1972a: Extension of covariance selection mathematics. Annals of Human Genetics 35(4):485–490.

———. 1972b: Fisher's 'Fundamental Theorem' made clear. Ann Hum Genet 36(2):129–40.

———. 1995: The nature of selection. Journal of Theoretical Biology 175(3):389 – 396.

Queller, D. C. 1992: A general model of kin selection. Evolution 46:376–380.

Rivoire, O. and Leibler, S. 2011: The value of information for populations in varying environments. J Stat Phys 142:1124–1166.

Schneider, T. D. 1997: Information content of individual genetic sequences. J Theor Biol 189(4):427–41.

———. 2000: Evolution of biological information. Nucleic Acids Research 28(14):2794–2799.

Schneider, T. D. and Stephens, R. M. 1990: Sequence logos: a new way to display consensus sequences. Nucleic Acids Res 18(20):6097–100.

Schneider, T. D.; Stormo, G. D.; Gold, L.; and Ehrenfeucht, A. 1986: Information content of binding sites on nucleotide sequences. J. Mol. Biol. 188:415–431.

Shannon, C. 1948: A mathematical theory of communication. Bell System Technical Journal 27:379–423,623–656.

Strelioff, C. C.; Lenski, R. E.; and Ofria, C. 2010: Evolutionary dynamics, epistatic interactions, and biological information. J Theor Biol 266(4):584–94.

Szostak, J. W. 2003: Functional information: Molecular messages. Nature 423(6941):689.

Taft, R. and Mattick, J. 2003: Increasing biological complexity is positively correlated with the relative genome-wide expansion of non-protein-coding DNA sequences. Genome Biology 5:P1.

Taylor, S. F.; Tishby, N.; and Bialek, W. 2007: Information and fitness. ArXiv.org preprint arXiv:0712.4382.

van Veelen, M. 2005: On the use of the Price equation. Journal of Theoretical Biology 237(4):412 – 426.

Wade, M. 1985: Soft selection, hard selection, kin selection, and group selection. Am Nat 125:61–73.

Wheeler, J. A. 1990: Information, physics, quantum: The search for links. Complexity, Entropy, and the Physics of Information, (edited by W. H. Zurek). Addison-Wesley, Redwood City, CA.

Wilke, C. O. and Adami, C. 2002: The biology of digital organisms. Trends in Ecology & Evolution 17(11):528–532.

Wilke, C. O.; Wang, J. L.; Ofria, C.; Lenski, R. E.; and Adami, C. 2001: Evolution of digital organisms at high mutation rates leads to survival of the flattest. Nature 412(6844):331–333.

Wiser, M. J.; Ribeck, N.; and Lenski, R. E. 2013: Long-term dynamics of adaptation in asexual populations. Science 342(6164):1364–7.